Published in final edited form as:

Comput Phys Commun. 2010 November 1; 181(11): 1896–1905. doi:10.1016/j.cpc.2010.07.022.

Browndye: A Software Package for Brownian Dynamics

Gary A. Huber^a and J. Andrew McCammon^{a,c,d,b}

^aHoward Hughes Medical Institute, University of California San Diego, La Jolla, CA 92093-0365

^bCenter for Theoretical and Biological Physics, University of California San Diego, La Jolla, CA 92093-0365

^cDepartment of Chemistry and Biochemistry, University of California San Diego, La Jolla, CA 92093-0365

^dDepartment of Pharmacology, University of California San Diego, La Jolla, CA 92093-0365

Abstract

A new software package, Browndye, is presented for simulating the diffusional encounter of two large biological molecules. It can be used to estimate second-order rate constants and encounter probabilities, and to explore reaction trajectories. Browndye builds upon previous knowledge and algorithms from software packages such as UHBD, SDA, and Macrodox, while implementing algorithms that scale to larger systems.

Keywords

Brownian dynamics; protein association; software; diffusion; kinetics

1. Introduction

One of the main determinants of biological structure and function is the interaction of two or more molecules, especially protein molecules. Understanding the dynamics of these bi-molecular interactions is important for the understanding of such cellular structures as the cytoskeleton (actin and tubulin, for example), ribosomes, chromosomes, and polymerases, as well as processes such as cell signalling and cell motility [1,2]. Furthermore, the encounter stages of such reactions, which are often the rate-limiting steps, are diffusion-limited. Therefore, use of Brownian dynamics is appropriate for such systems (See Reference [3] for an excellent review). For several decades, Brownian dynamics has found use in polymer and peptide simulations [4,5,6] and simulations of enzyme-substrate reactions [7,8,9,10]. More recently, it has found use in studying protein-protein association reactions [11,12].

The Browndye software package consists of two simulation programs and about 25 auxiliary programs for processing data. It can compute the 2nd-order rate constant of the encounter of two rigid bodies moving according to Brownian dynamics (BD), and it can also compute the probabilities of the two bodies moving from one binding mode to another, as well as

© 2010 Elsevier B.V. All rights reserved.

Correspondence to: Gary A. Huber.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

outputting the molecules' trajectories. It has functionality very similar to the packages SDA [11] and MacroDox[13], and like those two other packages, is intended primarily for simulations of large biological molecules. Its current limitations arise mainly from the approximation of flexible molecules as rigid bodies and the approximate nature of the force computations between the molecules. It also draws ideas from an earlier package, UHBD [8].

2. Model Description

Given unlimited computational resources, one would want to include every atom in the solvent and the large molecules in molecular simulations. Unfortunately, such detailed models are still impractical for simulations of macromolecular encounters because of the large number of individual objects and the disparity in time scales. Therefore, common practice is to model the solvent using continuum models, along with a simplified description of the macromolecules. A detailed model of the solvent is replaced with two simplified models. First, the thermal motion and hydrodynamic drag are replaced by a suitable stochastic force on the macromolecules. Second, the ion concentration and dielectric properties of the water are used in a continuum-based models for computation of the electrostatic forces between the macromolecules [14,15]. A final assumption of the Brownian model is that the solvent damping is very large compared to inertial effects, turning the macromolecules' trajectories into a random walk with no defined linear or angular velocities. The general equation describing the motion is [4,16,17,18,¹⁹,20,21,22,3]:

$$d \begin{pmatrix} \mathbf{X}_1 \\ \phi_1 \\ \mathbf{X}_2 \\ \phi_2 \end{pmatrix} = \frac{dt}{k_b T} \mathbf{D} \cdot \begin{pmatrix} \mathbf{F}_1 \\ \mathbf{T}_1 \\ \mathbf{F}_2 \\ \mathbf{T}_2 \end{pmatrix} + \sqrt{2 dt} \mathbf{S} \cdot \mathbf{w} + \nabla \cdot \mathbf{D} dt \quad (1)$$

where the vectors $\mathbf{x}_1$, $\mathbf{x}_2$, ϕ_1 , and ϕ_2 are the positions and orientations of the molecules, dt is the time step size, the matrix $\mathbf{D}$ is a generalization of the diffusivity which can also depend on positions and orientations, and the matrix $\mathbf{S}$ is the matrix square root of $\mathbf{D}$ (often computed using the Cholesky decomposition). The vector $\mathbf{w}$ comprises 12 independent random Gaussian variables, each with zero mean and unit variance. Computing the terms in this equation and updating the molecules' states with it is the main computational task of BrownDye.

The model for the rigid-body molecule is an assemblage of possibly overlapping charged spheres, grouped into "residues". Each sphere and residue can also have a "type". Although a common use of this model is to equate each sphere with an atom, and, in the case of proteins, to equate each residue with an amino acid, one can generalize this model to other molecular models. For example, coarse-grained models of proteins often treat each amino acid residue as a sphere.

In computing an association rate constant, it is necessary to decide when a reaction has taken place. A reaction criterion is specified by pairs of spheres, with each sphere on a different molecule. Each pair is also assigned a distance. The reaction criterion itself is further characterized by a number n . A pair is said to be "satisfied" if its sphere centers are closer than the assigned distance. If n pairs are satisfied, then the reaction is said to occur. For example, a criterion might be composed of all pairs of atoms forming hydrogen bonds in a protein-protein complex, and if any three of the pairs are within 5.1 Å. In addition to containing pairs, a reaction criterion can contain other reaction criteria; these inner criteria are treated like pairs to be satisfied. An example would be an enzyme with two active sites.

The binding of the substrate to each site would be represented by a separate reaction criterion, but the two reaction criteria could be lumped into another criterion to represent the overall reaction.

Entire reaction pathways can be specified by named **states** and named **reactions**. As an example, suppose the system comprises a substrate that can react first at one active site on an enzyme, and then can react at a second active site. When the molecules are initialized for a trajectory, the system is considered to be in the *reactants* state, and if they react, the system moves to the *product-1* state. The reaction itself is called *reaction-1*. The first reaction can then be followed by binding to the other site. The following state is called *product-2*, and the reaction is called *reaction-2*. An example of the input file specifying such a system is given in the documentation. The output of the Browndye software gives rate constants or probabilities for each named reaction.

The forces appearing in Eq. 1 fall into two categories: electrostatic and short-ranged. The electrostatic forces are computed from pre-computed solutions to the Poisson-Boltzmann equation describing the electric field around each molecule in isolation

$$\nabla \cdot [\epsilon \nabla \Phi] = -\rho - \lambda \sum_i c_i z_i \exp \left[-\frac{z_i \Phi}{k_b T} \right] \quad (2)$$

where Φ is the electric potential, ϵ is the dielectric of the material (either solvent or macromolecule), ρ is the permanent charge density in the macromolecule, and c_i and z_i are the bulk concentration and charge of the i^{th} ion species in the solvent. The factor λ depends on position and describes where mobile ions can penetrate; it typically takes a value of 1 in the solvent and 0 in the macromolecular interior. The electric potential around each molecule is typically computed by using the APBS software package [15] and stored in one or more rectilinear grids. Given the electric field around one molecule, Coulombic forces can be computed from its interaction with the sphere charges on the other molecule. This model alone is not sufficient, because as the molecules approach, their electric fields cause a build-up of induced charge at the dielectric boundaries between the macromolecules and solvent. This effect is partially included by computing so-called *desolvation forces* described below.

The short-ranged forces are of two types. The default short-ranged force is simply a hard-sphere exclusion force; if two spheres on different molecules overlap after a time step, the step is rejected. Browndye also includes an option to use Lennard-Jones forces between spheres; this is described by the pairwise potential energy between two spheres:

$$V(r) = \epsilon \left[\left(\frac{r_m}{r} \right)^{12} - 2 \left(\frac{r_m}{r} \right)^6 \right] \quad (3)$$

where r is the intersphere distance, ϵ is energy well depth, and r_m is the sum of the two radii. (This formula differs slightly from that seen commonly in the literature, in that the energy minimum occurs at the radii sum). The radius $r_{m,i}$ and energy ϵ_i are specified by the user for each sphere and residue type, with ϵ for a sphere pair computed from the mixing rule

$$\epsilon = \sqrt{\epsilon_i \epsilon_j} \quad (4)$$

It is possible to have both hard-sphere and Lennard-Jones forces in the system. Finally, Browndye includes an option for a “softer” Lennard-Jones-type interaction, the “8–6” force [23]:

$$V(r)=4\epsilon\left[\frac{3}{4}\left(\frac{r_m}{r}\right)^8-\left(\frac{r_m}{r}\right)^6\right] \quad (5)$$

Finally, the rate constant must be extracted from trajectories following Eq. 1 and the reaction criteria. This is done by running many trajectories, usually several thousand or more. A trajectory is started by positioning the molecule centers a distance b from each other, with randomly-chosen orientations. The molecules are stepped forward until they either react, in which case the trajectory ends, or they attain a distance q slightly larger than b . A random number is then used to decide whether the molecules escape from each other, thereby ending the trajectory. If they do not escape, they are again placed a distance b from each other, but with the orientations drawn from a specific probability distribution, and the trajectory continues. After the desired number of completed trajectories has been generated, the rate constant is computed from the ratio of reacted versus escaped trajectories [24]. In the case of complex reaction pathways, the movement of the molecules from state to state is recorded, until either the molecules escape or a final state with no reactions leading out of it is reached. Probabilities and rate constants are reported with 95% confidence intervals. The more trajectories are run, the more narrow are the intervals. An alternative method for computing the rate constant is also provided by the Weighted-Ensemble Brownian Dynamics method [25]; this can be useful in cases where the probability of reaction is very small compared to that of escape.

3. Software Description

The inputs into Browndye are the following. First of all, files describing the spheres' relative positions, radii, and charges, as well as their types and grouping into residues are required for each molecule. These files are in the PQR format, which is a variation on the Protein Data Bank file format. Next, files containing the electric field around each molecule are required; as mentioned above, these can be obtained from APBS and are in the DX format [26]. The reaction criteria are provided using an XML file which describes the pairs of spheres, their necessary distances, and number required for each reaction. Finally, an input file, also in XML format, is included, which gives information on the previously-mentioned files, physical properties of the solvent, and details regarding the simulation itself, such as various time-step size and force parameters, number of trajectories, information for outputting trajectory information, as well as many others.

The inputs are not fed directly into the simulation programs, but are processed by several auxiliary programs. This is done for several reasons. First of all, the preprocessing is done only once, but can still take some time; if it gets interrupted, it is much easier to restart and continue if the intermediate results are stored in files. For the software developer, it is much easier to maintain and change a collection of programs than one large, “black-box” program. Perhaps most importantly, the user can also replace any of the intermediate files with his own information and continue the preprocessing; this gives considerable flexibility. The user can also run any of the auxiliary programs on his data. For example, several users have opted to use a collection of “effective charges” (described below) which is more robust than what is currently provided by Browndye.

With the exception of the DX files containing the electric field data, the intermediate files are in XML format; even the PQR files are converted to an equivalent XML format. Using

XML files has two main advantages. For the user, the tags provide a human-readable commentary on the data. For the programmer, changing the format of the file by adding or removing information in the form of XML tags can be done without being required to rewrite the software that parses the file. It is also possible to automatically check an XML file against an expected format and to quickly flag any inconsistencies, although Browndye does not currently make use of this feature. The programs of Browndye use the freely available and light-weight XML parser Expat [27].

Almost all of the auxiliary programs are written in the Objective Caml (or OCaml) language [28], which is a variant of the ML computer language. OCaml is a statically-typed, garbage-collected functional language which combines the conciseness, safety, and ease of programming offered by functional and scripting languages with the performance of procedural languages such as C++ and Fortran. Each auxiliary program is small enough that it can be treated as a black box, so the general unfamiliarity of scientists with OCaml should not be an issue.

On the other hand, the two simulation programs **nam_simulation** and **we_simulation**, which perform the trajectory and weighted-ensemble simulations, are written in C++. First of all, Ocaml might not be available on the computers at the various supercomputer centers. While the user can almost always preprocess the data on his own computer where he has some control, he might want to run the actual time-consuming simulation elsewhere. Also, a future goal of this project is to provide various generic algorithmic building blocks for building other software packages, and the templating facility of C++ is ideal for this. Finally, at this time, the Ocaml language does not offer support for concurrent multi-threaded execution on several processors.

Because the execution of Browndye consists of a moderately complex network of files and programs, another program, called **bd_top**, is used to keep track of which files are up-to-date and which auxiliary programs still need to run, and calls the necessary programs with appropriate arguments. This typically is the first program in the Browndye suite to be run, and it takes the XML input file mentioned above as its input. The program **bd_top** is written using an Ocaml module called the **orchestrator**, which is similar to the “make” utility on many Unix-based systems [29]. Like the “make” utility, it can compare the last-updated times of files and decide which programs to run. In addition, it can also compare tagged quantities between two XML files and extract arguments to programs from XML files. As Browndye grows and evolves, this allows the **bd_top** program to be updated accordingly. After the program **bd_top** is run, which in turn generates many more XML and DX files as well as an input file into one of the simulation programs, one of the simulation programs is run (Fig.1). Because of the computationally intensive nature of the simulations, the two simulation programs have been written to make use of systems with several processors and shared memory. They make use of *pthreads*, the standard interface for systems-level multithreaded programming [30]. In the single-trajectory **nam_simulation** program, each trajectory can be run independently, so a thread is created for each processor. When a thread finishes with a trajectory, it pulls another future trajectory off of a queue, thus ensuring that each thread is fully occupied until the very end of the simulation. In the weighted-ensemble **we_simulation** program, all of the two-molecule system copies are advanced simultaneously at each time step, and the system time step advancements are distributed among the threads using the same queuing method. The simulation programs are structured so that most of the data associated with the molecules is shared among the threads. In particular, the data grids, which hold the electrostatic potentials and other quantities, use the bulk of the total memory required. Fortunately, each large data grid can be read by any of the threads; no grid needs to be duplicated.

At present, there are two types of outputs from the simulation programs. As they run, they update an XML file containing raw information on which trajectories have escaped or completed given reactions. This information, at any time, may be given to another program, **compute_rate_constant**, which then computes second-order rate constants, along with their confidence intervals, and the relative probabilities of following the various reaction pathways. The other type of output, from `nam_simulation`, consists of the simulation trajectories themselves and their states. The output state consists of the relative positions and orientations, various components of the interaction energy, and the current reaction state. Because these files can be quite large, the numerical data for twenty adjacent states is encoded into a MIME format and embedded in XML in order to reduce the number of memory-consuming XML tags. In addition to trajectory files, an “index” file is output, which allows subsequent post-processing programs to quickly find certain trajectories with desired properties. The main post-processing program, **process_trajectories**, can be used to extract out portions of trajectories that correspond to desired reaction paths. The resulting sub-trajectories can be processed further by the program **xyz_trajectory** to generate files for visualization with the molecular visualization package VMD [31].

4. Algorithms and Optimizations

Although Browndye has functionality similar to other packages mentioned above, it includes several algorithms and optimizations to help its performance scale well to larger systems. Most of these techniques are concerned with computing long-ranged forces, resolving collisions, and propagating the trajectories. Because each molecule is made up of a large number of spheres (N), any algorithm whose execution time would scale as N^2 is avoided, while algorithms that scale as N or $\log(N)$ are included, even if their benefit is not realized as strongly on smaller systems. Algorithms that help to reduce the number of distinct steps, along with their required force calculations, are also included.

As mentioned above, most of the required memory is consumed by the field grids; these contain the electric field around the larger molecule, and desolvation fields (see below) around both molecules. For points near the surface, a higher resolution is required, which in turn requires smaller spacing between the grid points. Further away, however, a coarser grid can be used. Browndye allows the use of such nested grids in order to save memory; any number of such grids can be imported from the output of APBS. Outside of the outermost grid, the electric potential data in the outer points is fit, using least-squares, to a Cartesian multipole expansion [32], out to quadrupole level, of the screened Coulombic interaction. Any point outside of the grids that needs a value of the potential or its gradient uses the multipole expression.

Perhaps the most prominent force is the Coulombic interaction among the two molecules; this can be computed from the charge at each sphere on the smaller molecule and the gradient of the electric potential at that point, usually by trilinear interpolation from the nearest eight grid points. Often, the “smaller molecule” can itself be quite large, and the charges on the large molecule are also used in the desolvation forces below. Looping over large numbers of charges at each time step can be time-consuming, so several approaches are used to reduce the effort.

The first step is to reduce the number of charges to a smaller set of “effective charges”. One commonly-used approach, used in the SDA package, is to fix a charge position on each charged residue and to perform a least-squares fit of the charges to the surrounding electric field [33]. The main drawbacks are that the effort to compute the charge values scales unfavorably with large systems, and a certain amount of experimenting with a regularization parameter is required in order to avoid unphysical charge values. Currently, Browndye uses

a simpler, “test charge” approach by placing one point charge at the center-of-charge at each charged residue, with the point charge value equal to the total residue charge. As mentioned above, however, it is possible and desirable to import effective charges from SDA or other methods by replacing one of the intermediate XML files and running `bd_top` again to update the other files. Future improvements to Browndye will include more robust and quickly-computed effective charges.

Given a set of effective charges, it is possible to lump point charges together dynamically as the simulation proceeds. If the electric field is relatively smooth in the region of a set of charges, a coarse-grained description can be used in order to avoid looping over all of the charges and computing the field gradient at each point. In general, the further away from a charge source one is, the smoother is the field, so a decision to use a coarse-grained model can be based on distance from the other molecule during the simulation.

Browndye surrounds the molecule with a box, and divides the box into two smaller boxes along its longest axis so that equal numbers of charges fall into the two child boxes. This process is recursively repeated until a small number of charges remains in each bottom-level box. Each box has 64 quadrature points at which the electric potential is evaluated, and the obtained values are used to compute the force and torque on the box (Appendix). The model used, which is based on Chebyshev interpolation of the potential, assumes that the field is fairly smooth on the scale of the box. If the field is not smooth, then the child boxes need to be used for the force and torque evaluation. At each time step, the algorithm starts at the top box and recurses down the hierarchy of boxes. If the distance from the box center to the nearest point charge on the other molecule becomes less than a certain factor times the diagonal length of the box, then the children are used. An efficient algorithm, described below, is used to find the distance. If bottom-level box is reached, then the force and torque are computed from the point charges themselves and the field gradient.

The desolvation penalty force is computed by using a formula very similar to the one used in the SDA package [11], but with a prefactor inspired by the generalized Born model[34]. The contribution due to one effective charge interacting with the other molecule is:

$$\varepsilon = \frac{q^2 \alpha (\varepsilon_p^{-1} - \varepsilon_s^{-1})}{32\pi^2 \varepsilon_0} \int (1 + \kappa r)^2 \exp(-2\kappa r) r^{-4} dr \quad (6)$$

where ε_0 is the vacuum permittivity, ε_p and ε_s are the dimensionless dielectrics of the molecule interior and the solvent, κ is the inverse Debye length of the solvent, q is the value of the effective charge, r is the distance from the charge, and the integral is over the interior of the other molecule. The factor α , whose default is 1, may be included for a better match to the energy. The factor without the q^2 term is precomputed for each molecule and is stored on a grid, with r corresponding to the exterior grid points. The integral is performed using a Cartesian-based fast multipole method [35], allowing this preprocessing step to scale linearly with the size of the molecule. The desolvation grids match the shape and number of the electric potential grids. The squared effective charges are lumped in a recursive box structure as described above, and during the simulation, they interact with the precomputed desolvation in the manner as the charges interact with the electric field. The interior of the molecule is represented by a grid of integers (1 for inside, 0 for outside) of the same geometry as the smallest electrostatic grid enclosing the whole molecule. The generation of this grid is described below.

Hydrodynamic interactions are handled by using the variation of the Rotne-Prager tensor [36] developed by Garcia de la Torre and Bloomfield [37] for two spheres:

$$\mathbf{D}_{ij} = \frac{kT}{6\pi\mu\sigma_i} \mathbf{I} \delta_{ij} + \frac{kT}{6\pi\mu r_{ij}} \left[\left(\mathbf{I} + \frac{\mathbf{r}_{ij}\mathbf{r}_{ij}}{r_{ij}^2} \right) + \frac{\sigma_i^2 + \sigma_j^2}{r_{ij}^2} \left(\frac{1}{3} \mathbf{I} - \frac{\mathbf{r}_{ij}\mathbf{r}_{ij}}{r_{ij}^2} \right) \right] (1 - \delta_{ij}) \quad (7)$$

where μ is the solvent viscosity, $\mathbf{r}_{ij}$ is the vector from sphere i to sphere j , and σ_i is the radius of sphere i . The diffusion matrix in Eq. 1 is given by

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}_{11} & 0 & \mathbf{D}_{12} & 0 \\ 0 & \mathbf{D}_{r1} & 0 & 0 \\ \mathbf{D}_{21} & 0 & \mathbf{D}_{22} & 0 \\ 0 & 0 & 0 & \mathbf{D}_{r2} \end{pmatrix} \quad (8)$$

where the rotation matrices are given by [20]

$$D_{ri} = \frac{kT}{8\pi\mu\sigma_i^3} \quad (9)$$

This model neglects the hydrodynamic rotation-rotation and rotation-translation coupling. One significant advantage of the Rotne-Prager model is that the divergence of the diffusion matrix in Eq. 1 is zero.

Because the molecules, in general, are not perfect spheres, their effective radii are computed using the Monte Carlo surface integral method of Hansen [38]. One of the intermediate files is a grid of points with an indicator of whether the point is inside or outside the molecule (see below). Pieces of the molecule's surface, for this purpose, are defined within each cube whose vertices are made up of both inside and outside points. A plane is assumed to cut through each edge shared by an inside and outside point, and the areas and centers for each bit of surface are computed for each surface cube.

The effective gap between the molecules is computed by assuming that the molecules are ellipsoids. To find the equivalent ellipsoid of a molecule, the moment of inertia matrix is computed for the molecule interior. This matrix is inverted to obtain a transform that maps the points of the molecule onto a solid with an isotropic moment matrix. Then, the smallest enclosing sphere (see below) of the transformed molecule subspheres is found. This sphere is deformed back into an ellipsoid using the moment matrix. During the simulation, the intersection of the points on the two ellipsoids with the line segment connecting their centers is computed, and the gap is approximated as the distance between the two intersection points. The distance r_{12} in Eq. 7 is the sum of the effective radii and the effective gap.

One computational challenge in BD simulations is figuring out when a collision takes place between two rigid assemblages of spheres. One method would be to compute the distance between every pair of spheres, but the required time would scale unfavorably for large systems. The SDA and MacroDox packages address this by pre-computing an exclusion grid for one molecule, and pretending that all the spheres on the other molecule have the same size. The solution used by Browndye is to enclose each molecule in a hierarchy of spheres which enclose the various parts of the molecules, and test the pairs of fewer and larger spheres, recursing downward if necessary. This is a simple version from the class of Hierarchical Bounding Volume methods [39]. First, a sphere is generated that closely approximates the minimal enclosing sphere of the molecule's spheres. The minimal enclosing sphere is approximated by first computing the minimal enclosing sphere of the

centers using the method of Welzl [40]. Then, the sphere radius is increased until all of the spheres are inside. After this, the moment of inertia matrix is computed for the sphere centers, and sphere set is divided into equal halves by a plane perpendicular to the principal axis corresponding to the largest eigen-value. Enclosing spheres are constructed for the sphere subgroups, which are divided again, until just a small number of molecule spheres remain in the group. During the simulation, the outer-most spheres are tested for over-lap. If there is no overlap, then there is no collision. Otherwise, the two spheres on each molecule which are one level down are tested for collisions with each other (four different tests), and the algorithm recurses down until two atom-spheres are found which collide.

In effect, Browndye keeps the large molecule fixed, while letting the smaller molecule move. Of course, both molecules move, but after each move, the system is translated and rotated so that the larger molecule is back in its original position and orientation. So, the positions of the atom-spheres of the larger molecule are readily available. The situation is more complex with the smaller molecule; the positions of the enclosing spheres and the atom-spheres must be computed by transforming from the original positions. Because not all spheres are tested, it does not make sense to transform all of them, so only those that are being tested should be transformed. So, the first time a sphere is tested, its transformed position is computed and stored, since the sphere may be tested several times. When the transformed position is stored, a flag is set to indicate that this has happened and to keep the work from being duplicated later. After the collision testing is complete, the algorithm recursively descends from the top sphere and clears the flags for the next test. In using multiple threads, then, the flags and the transformed positions of the spheres of the smaller molecule must be duplicated for each thread, while all of the threads can use the same copy of the positions of the larger molecule.

The short-ranged Lennard-Jones forces are computed using the same idea. An effective radius is computed for each sphere, with the edge of the sphere marking the point where the potential energy is negligible. During the downward recursion, all in-range interactions are computed and added to the total force and torque. The main difference between this and the collision test is that all in-range pairs are tested and processed, rather than stopping once a collision is found.

A related algorithm, which is used in several places in Browndye, is the determination of the nearest atom-sphere to a given point outside of the sphere collection. A similar class of algorithms has been used for ray-tracing applications in computer graphics.[41] First, the distance between the point and the surface of the top sphere is computed, and the top sphere is put on a priority queue. This priority queue is ordered by distance from the point, so the item that is pulled off the queue is the closest sphere to the point. The algorithm proceeds by pulling a sphere off the queue, and checking to see if the nearest possible atom-sphere is greater than an upper bound. If it is, then the sphere is discarded; otherwise, its two child spheres are enqueued. If the distance from the test point to the furthest point in the dequeued sphere is less than the upper bound, the upper bound is set to that distance. This prevents the need to check spheres that are known with certainty not to contain a nearest atom-sphere. Spheres are pulled off the queue until one atom-sphere remains that satisfies the upper bound. Because the priority queue is implemented using the heap data structure, adding and pulling off items can be done very efficiently.

Some computer time can be saved by eliminating ahead of time those spheres that will not participate in a collision. First, a set of surface spheres are defined by rolling a probe ball across the surface using the algorithm of Connolly.[42,43]. The neighbors of a surface sphere are those which are touched at the same time as itself by the probe sphere, and the neighboring sphere centers define a closed mesh of triangles. All of the spheres that have not

been touched by the probe sphere are tested to see if their centers are inside the mesh or any of the surface spheres. This is done by picking a point well outside the mesh, and computing the number of intersections between a line connecting the two points and the mesh triangles (more details below). Those that are inside are discarded, and the ones not inside are called “danglers”. The outside spheres are removed and saved elsewhere, while the probe sphere is rolled across the set of danglers. Since the set of danglers is usually divided into isolated clusters and the probe sphere can roll across only one cluster, the process is repeated until the clusters have been processed and divided into inside and outside spheres, and no more spheres remain.

Once the closed meshes of triangles and set of dangling outside spheres are generated, the grid of interior points is generated. A point is considered to be “inside” if it is overlapped by one of the spheres expanded by an additional exclusion distance, or if it is inside one of the closed triangular meshes generated above. Because there is a large number of grid points and often a large number of spheres and triangles, it is important to have an efficient algorithm that avoids testing every pair of grid point and sphere or triangle. The spheres are grouped into a tree-of-spheres data structure as used for collision detection above, and the triangles are likewise grouped into such a data structure. In order to test a point, the nearest sphere is found; if it overlaps, then we know that it is inside. Otherwise, another point is selected which is distant enough from the grid center to guarantee that it is outside the molecule.

We then have to determine how many triangles intersect the line segment connecting the test point with the outer point. If the number of intersections is odd, then the test point is inside; otherwise, it is outside. This algorithm proceeds recursively by finding the distance from the line segment to the outermost sphere enclosing the triangles. If the line segment does not intersect the sphere, then there are no intersections. Otherwise, the two child spheres are likewise tested, with no further work required for any non-intersecting spheres. Eventually, individual triangles are tested, and the total number of intersections are added up. This method eliminates the need to test most of the triangles. In the rare case that a point of intersection is on an edge between two triangles, the outside point is perturbed very slightly and another attempt is made.

Another gain in efficiency can be made by using the fact that if an outside point is a more than certain distance from the surface of the nearest sphere, then all other points within that distance of the first point are also outside. Conversely, if an interior point is known to be a certain distance from the surface of the nearest triangle, all other points within that distance of the first point are also inside. So, rather than stepping through the grid points sequentially, a self-avoiding sequence of grid points is generated by incrementing an integer index, reversing the bits, and using every third bit, starting with the first bit, to generate the index in the x-direction. The y-index and z-index are likewise generated by starting at the second and third bits. If the status (inside or outside) of the selected grid point is unknown, it is tested, and the distance to the nearest sphere or triangle is found. Then, all the points within that distance are set to have the same status. In this way, large chunks of points are determined without having to perform a test for each one.

Many of the existing packages for BD use the Northrup-Allison-McCammon (NAM) algorithm [44]. This algorithm starts the smaller molecule on the b -radius described above, and steps the system forward until a final reaction has been reached, or a larger q -radius is reached. In the traditional algorithm, the trajectory is terminated upon reaching the q -radius, and the q -radius must be considerably larger, in many cases, than the b -radius. Also, the required q -radius is difficult to compute, leading most users to use a large value in order to be safe. A newer version by Luty, McCammon, and Zhou(LMZ)[24] allows one to use a

smaller q -radius, but then either places the smaller molecule back on the b -sphere according to an analytically computed probability distribution without yet terminating the trajectory, or assuming an escape and terminating the trajectory (Fig. 2). The original Luty-Zhou-McCammon algorithm does not take into account the change in orientation of the smaller molecule during its presumed journey from the q -sphere back to the b -sphere. The implementation in Browndye does take this into account, using an enhanced probability distribution that includes not only the new starting position but also the time taken to get there. Given a travel time, the algorithm then generates a change in orientation from a probability distribution, using a series formula that is precomputed and stored (unpublished results).

When using the weighted-ensemble method with the `we_simulation` program, the system copies (molecule pairs) are started with a separation of b , and the probabilities of reaction versus escape are computed from the relative probability fluxes. As the preprocessing step, the bins are generated by starting the copies with a separation of q , and running the copies until they react, setting up the bin partitions as described in Huber and Kim [25]. During the bin generation phase, the system copies “bounce back” if they go past a separation of q . Rather than using the simple bootstrap method as in the original paper, the confidence intervals on the probabilities are derived using the Stationary Bootstrap method for time series [45].

Although there are no restrictions on the size of the q -sphere in the LMZ algorithm, the b -radius must be large enough so that the force between the molecules depends only on the center-to-center distance of the molecules and not their orientation. Browndye can automatically determine the smallest radius that satisfies this requirement. One of the auxilliary programs first determines the smallest sphere about each molecule where the electric field is uniform, within a certain tolerance, on the surface of the sphere; these radii are r_{f1} and r_{f2} . Also, the maximum distance from the center among the point charges is computed for each molecule (r_{c1} and r_{c2}). The b -radius is taken to be

$$b = \max(r_{f1} + r_{c2}, r_{f2} + r_{c1}) \quad (10)$$

which ensures that each charge on one molecule interacts only with a radially symmetric field from the other molecule. Currently, the q -radius is taken to be 1.1 times the b -radius.

One important aspect of BD simulations is the selection of the time step size. Not only can the choice of time step size significantly affect performance, the computed reaction rates can be quite sensitive to the size if it is too large. Browndye attempts to determine an appropriate time step based on the geometry of the two molecules, and then corrects and refines the time step, if necessary, if the change in the forces and torques are too large. The following expression is used:

$$r_{mid} = \frac{1}{2}(q+b) \quad (11)$$

$$g_o = q - r + \mathcal{H}(r - r_{mid})(r - r_{mid}) \quad (12)$$

$$\mathcal{D}=(\mathbf{D}_{11}+\mathbf{D}_{22}) \cdot \frac{\mathbf{r}}{r} \quad (13)$$

$$\Delta t_o = \frac{\alpha_o g_o^2}{\mathcal{D}} \quad (14)$$

$$\Delta t_g = \frac{\alpha_g g^2}{\mathcal{D}} \quad (15)$$

$$\Delta t_r = \frac{\alpha_r \rho_{min}^2}{\mathcal{D}} \quad (16)$$

$$\mathcal{D}_r = \max(\mathcal{D}_{r1}, \mathcal{D}_{r2}) \quad (17)$$

$$\Delta t_{max} = \frac{\pi^2}{\mathcal{D}_r} \quad (18)$$

$$\Delta t_{min} = \frac{\Delta x^2}{2\mathcal{D}} \quad (19)$$

$$\Delta t_{r,min} = \frac{\Delta x_r^2}{2\mathcal{D}} \quad (20)$$

$$\Delta t = \min(\Delta t_{mat}, \min(\max(\Delta t_{min}, \min(\Delta t_g, \Delta t_o)), \quad (21)$$

$$\max(\Delta t_{r,min}, \Delta t_r))) \quad (22)$$

where q and b are the q - and b -radii, r is the center-to-center distance, $\mathcal{H}$ is the Heaviside step function, α_o , α_g , and α_r are various tolerances, g is the effective gap, $\mathcal{D}$ is the relative diffusivity of the two molecules along the axis through both centers, $\mathcal{D}_{r1}$ and $\mathcal{D}_{r2}$ are the rotational diffusivities of the two molecules, ρ_{min} is the minimum of all reaction coordinates that are relevant for the current state, Δx_{min} is an average minimum distance per time step. The parameter $\Delta x_{r,min}$ is like Δx_{min} but can be made even smaller; this allows very small time steps near reactive configurations. The various time step components have physical interpretations. The time step Δt_o makes sure that too big of a time step is not taken near the

q -sphere. The time step Δt_g ensures that small steps are taken when the molecules collide, and Δt_r ensures that small steps are taken near a reactive configuration. The time step Δt_{max} ensures that the rotations are small enough (see below), and the steps Δt_{min} and $\Delta t_{r,min}$ put limits on how small the time steps can become. The dimensionless α -parameters and the parameters Δx_{min} and $\Delta x_{r,min}$ are given reasonable defaults which can be overridden by the user.

In addition to using the geometrically determined time step Δt , Browndye bounds the time steps based on changes in $\mathcal{D}$ and in the force and torques between the molecules. Eq. 1 makes the assumption that these quantities do not change much over the course of a time step, so if that assumption is violated, the time step must be reduced. One simple scheme would be to reject the time step and just take another, smaller one. This, however, introduces a bias, so Browndye uses a more elaborate scheme. One can imagine a random walk with some vector $\mathbf{W}$ driven by the Gaussian variables $\mathbf{w}$ in Eq. 1:

$$\mathbf{W}_i = \sum_{j=1}^i \sqrt{2dt} \mathbf{w}_j \quad (23)$$

Suppose we took too large a time step. We would still want to keep the random walk up to this point to avoid biasing the results, but we would subdivide it by adding an extra $\mathbf{W}$ according to the formula

$$\mathbf{W}_{i-1/2} = \frac{1}{2}(\mathbf{W}_{i-1} + \mathbf{W}_i) + \sqrt{\frac{\Delta t}{2}} \mathbf{N} \quad (24)$$

where $\mathbf{N}$ is a vector of independent Gaussian variables with zero mean and unit variance. Then, the $\mathbf{w}$ variables from this refined random walk are

$$\mathbf{w}_{i-1/2} = \mathbf{W}_{i-1/2} - \mathbf{W}_{i-1} \quad (25)$$

$$\mathbf{w}_i = \mathbf{W}_i - \mathbf{W}_{i-1/2} \quad (26)$$

and both are used with half the original time step. If, upon using the reduced time step, the conditions are still violated, the algorithm backs up, and subdivides the random walk again. All values of $\mathbf{W}$ in the future that have been calculated from subdividing the random walk are used, and repeated subdivisions can lead to a chain of $\mathbf{W}$'s that must be followed. Eventually, the algorithm reaches the end of the chain of $\mathbf{W}$'s, and the values of $\mathbf{w}$ are generated from a Gaussian distribution again. The time step must be subdivided if any of the following inequalities is satisfied, for either molecule:

$$|\mathbf{F}_i - \mathbf{F}_{i-1}| |\mathbf{x}_i - \mathbf{x}_{i-1}| > 0.02kT \quad (27)$$

$$|\mathbf{T}_i - \mathbf{T}_{i-1}| |\phi| > 0.02kT \quad (28)$$

where ϕ is the rotation angle between the two successive orientations. In addition, the subdivision is triggered if there is a collision, or if this inequality is satisfied:

$$|\mathcal{D}_i - \mathcal{D}_{i-1}| > 0.02|\mathcal{D}_{i-1}| \quad (29)$$

One difficulty that is sometimes encountered with collisions between groups of hard spheres is that they can get stuck, and no reasonable number of small time steps can get them apart. Browndye detects this occurrence by keeping track of how many collisions have occurred since the last successful step. If a certain number (currently 10) of collisions occur in a row, then the two molecules are moved apart, along the line connecting their centers, by amount of overlap between the colliding atom-spheres. If there is still a collision, then the molecules are further moved apart in intervals of Δx_{min} until there is no more collision.

Another enhancement removes the need to take small steps near the absorbing q -sphere by using the method of Lamm and Schulten [46] to solve the unsteady-state diffusion equation near the boundary. Although the boundary is curved, the center of the smaller molecule is close enough so that the boundary can be considered flat, resulting in a one-dimensional diffusion equation. When the center of the smaller molecule is outside the b -sphere and inside the q -sphere, and is closer to the q -sphere, the molecules are stepped forward in time as usual. If their separation goes beyond the q -radius, then the smaller molecule either escapes or is put back onto the q -sphere, as described above. If not, then a survival probability is computed:

$$P_{sur} = 1 - \exp\left(-\frac{(q - r_{old})(q - r_{new})}{\mathcal{D}dt}\right) \quad (30)$$

where r_{old} and r_{new} are the old and new separations. A uniform random number between 0 and 1 is generated; if it falls above P_{sur} , the smaller molecule is assumed to be absorbed by the q -sphere and is processed as described above. This additional treatment takes into account the possibility that, during a large time step, the separation might become greater than q , even if the final separation is not. Without it, Browndye would have to take very small time steps in order not to underestimate the absorption at the q -sphere.

When using Eq. 1 to step the system forward, adding together the results of the changes in orientation (ϕ_1 and ϕ_2) is not as straightforward as it is for the positions. Numerically, it is most convenient to convert the incremental rotations, as well as the rotation matrix of the smaller molecule, and combine the rotations using quaternion multiplication. The resulting quaternion is then converted back to a rotation matrix. The time-stepping algorithm above includes a provision to limit the size of the rotations, since the rotational part of Eq. 1 is not valid for large changes in orientation.

5. Limitations

Currently, Browndye's model has many limitations. The molecules are treated as rigid bodies, ruling out realistic simulations of systems where molecular flexibility is significant for function. Inclusion of flexibility is currently in progress. The desolvation energy term is essentially an approximation of an approximation, and use must be made of the fudge factor in Eq. 6. The reaction criteria, especially the distances, can be used as another adjustable parameter, although one might conjecture that, given a good interaction model, the reaction rates would not be very sensitive to the reaction distances. There are no built-in hydrophobic forces, although one can approximate them using the Lennard-Jones forces. Nevertheless, Browndye is potentially useful for computing relative rates among different experimental conditions and mutants, and for exploring the trajectories to gain insight into receptor-ligand encounters. Future work on the model will include a better desolvation and effective charges

model, flexibility and mobile side chains. Future work on the software itself will include modularization and documentation of the components to enable other developers to use the above algorithms in their own software.

6. Example

An example of an diffusion-limited, electrostatically-driven protein-protein interaction is the binding of thrombin and thrombomodulin, two proteins in the blood clotting cascade. After the PDB structures of the two molecules were suitably preprocessed (e.g., hydrogens added), the electric fields around each were generated using APBS, assuming salt concentration of 0.15 mM. The pairs of spheres defining the reaction were taken from possible hydrogen-bonding pairs. Using a criterion of three required pairs at a separation of 4.45 Å, and a million trajectories, a rate of $7.0 \times 10^6 \pm 0.5 \times 10^6 M^{-1}s^{-1}$ was obtained, closely matching the experimental rate constant of 6.7×10^6 [47]. Using the same conditions but changing the salt concentration to 0.1 mM gave the rate constant $12.0 \times 10^6 \pm 0.8 \times 10^6 M^{-1}s^{-1}$, which is within the experimental error bars of the experimental value of $15.1 \times 10^6 M^{-1}s^{-1}$. Each run, which took place on 4 Intel Xeon processors, took about a day. (This is the system used in the example used in the software distribution, except that the reaction criteria in the example are changed to unrealistically large values in order to allow the user to compute “rates” after just a few minutes.)

7. Availability

The Browndye software package is freely available at the website browndye.ucsd.edu. The package and its components are distributed under an MIT-style license, and there are no components that are restricted by a GNU license. It runs under Linux, and can run on both 32-bit and 64-bit processors.

Acknowledgments

We are grateful to Maciej Dlugosz for patiently enduring a first draft of this software and providing valuable feedback. We also thank Barry Grant for being the first user of many of Browndye’s features and providing more valuable feedback. We thank Adam van Wynsbeghe for providing the preprocessed thrombin and thrombomodulin PDB files. This work has been supported by the National Institutes of Health, National Science Foundation, the Howard Hughes Medical Institute, the Center for Theoretical Biological Physics at UCSD, and the National Biomedical Computational Resource.

Appendix: Charge Lumping

A useful method for representing a function of one variable on the interval $[-1:1]$ is Chebyshev Approximation:

$$f(x) \approx \sum_{i=0}^{N-1} c_i T_i(x) \quad (31)$$

with coefficients

$$c_i = (1 - \frac{1}{2} \delta_{0i}) \sum_{m=1}^N T_i(x_m) f(x_m) \quad (32)$$

where T_i is the i th Chebyshev polynomial and the x_i are the N zeros of polynomial T_N , and δ_{0i} is the Kronecker delta. This is a useful approximation, because it tends to spread the error out over the interval.

The same method can be extended to three dimensions, where x, y , and z all range from -1 to 1 :

$$f(x, y, z) \approx \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} c_{ijk} T_i(x) T_j(y) T_k(z) \quad (33)$$

with coefficients

$$c_{ijk} = \frac{8}{N^3} \left(1 - \frac{1}{2} \delta_{0i}\right) \left(1 - \frac{1}{2} \delta_{0j}\right) \left(1 - \frac{1}{2} \delta_{0k}\right) \sum_{m=1}^N \sum_{n=1}^N \sum_{p=1}^N T_i(x_m) T_j(y_n) T_k(z_p) f(x_m, y_n, z_p) \quad (34)$$

For example, if one wanted to approximate a function to fourth order, corresponding to $N = 4$, one would evaluate the function at the 64 points of the Cartesian product of the zeros of T_4 and use that information to compute the coefficients.

Browndye approximates the electric potential using a 3-D Chebyshev interpolation, but using a different order of calculation and storage. Given a collection of point charges, a box is constructed about it. We scale and shift the positions so that the x, y , and z coordinates of the box range from -1 to 1 ; the hat on the positions, force, and torque denote the scaled quantities. Browndye uses $N=4$ as described above, resulting in the 64 points. If there are more than 64 point charges inside the box, we would like to evaluate the electric potential (in this case, interpolating from an APBS grid) at just the 64 points and obtain the total force and torque on the collection of charges. This, of course, can only be done if the potential field is smooth enough in that region to be accurately represented by the Chebyshev approximation.

Given an approximate potential V , the potential energy due to N_c charges is

$$E = \sum_{s=1}^{N_c} q_s V(\hat{\mathbf{x}}_s) \quad (35)$$

where q_s and $\hat{\mathbf{x}}_s$ are the s^{th} charge and position. Likewise, the force and torque are

$$\hat{\mathbf{F}} = - \sum_{s=1}^{N_c} q_s \nabla V(\hat{\mathbf{x}}_s) \quad (36)$$

$$\hat{\mathbf{T}} = - \sum_{s=1}^{N_c} q_s (\hat{\mathbf{x}}_s - \hat{\mathbf{X}}_0) \times \nabla V(\hat{\mathbf{x}}_s) \quad (37)$$

where $\hat{\mathbf{X}}_0$ is the point about which the torque is measured. From the preceding three equations, it can be shown that a linear relation exists between the potential values at the quadrature points on one hand, and the energy, force, and torque on the other:

$$E = \hat{\mathbf{K}}_E \cdot \mathbf{V} \quad (38)$$

$$\hat{\mathbf{F}} = \hat{\mathbf{K}}_F \cdot \mathbf{V} \quad (39)$$

$$\hat{\mathbf{M}} = \hat{\mathbf{K}}_M \cdot \mathbf{V} \quad (40)$$

where $\mathbf{V}$ is the vector of potential values, $\hat{\mathbf{K}}_E$ is a vector of length N^3 , and $\hat{\mathbf{K}}_F$ and $\hat{\mathbf{K}}_M$ are N^3 by 3 matrices. We define the coefficients

$$C_{mnp,ijk} = \frac{8}{N^3} (1 - \frac{1}{2}\delta_{0i})(1 - \frac{1}{2}\delta_{0j})(1 - \frac{1}{2}\delta_{0k}) T_i(\hat{x}_m) T_j(\hat{y}_n) T_k(\hat{z}_p) \quad (41)$$

and the Green's functions

$$V_{mnp}(\hat{x}, \hat{y}, \hat{z}) = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} T_i(\hat{x}) T_j(\hat{y}) T_k(\hat{z}) \quad (42)$$

The Green's function indexed by mnp is the result of plugging a one into the potential value at the quadrature point indexed by mnp and zero into the all the other values. The energy vector is then

$$\hat{K}_{mnp}^E = \sum_{s=1}^{N_c} q_s V_{mnp}(\hat{\mathbf{x}}_s) \quad (43)$$

where the quadrature points are indexes using mnp . The force and torque matrices are

$$\hat{K}_{r,mnp}^F = - \sum_{s=1}^{N_c} q_s \frac{\partial V_{mnp}(\hat{\mathbf{x}}_s)}{\partial \hat{x}_r} \quad (44)$$

$$\hat{K}_{r,mnp}^M = \sum_{s=1}^{N_c} q_s \sum_{i=1}^3 \sum_{j=1}^3 \varepsilon_{rij} \frac{\partial V_{mnp}(\hat{\mathbf{x}}_s)}{\partial \hat{x}_i} (\hat{x}_{s,j} - \hat{x}_{0,j}) \quad (45)$$

where ε_{rij} is the Levi-Civita symbol. These matrices can be constructed before the simulation.

The transformation from the scaled coordinates to the actual coordinates is

$$x = \frac{1}{2}(H_x + L_x) + \frac{1}{2}(H_x - L_x)\widehat{x} \quad (46)$$

$$y = \frac{1}{2}(H_y + L_y) + \frac{1}{2}(H_y - L_y)\widehat{y} \quad (47)$$

$$z = \frac{1}{2}(H_z + L_z) + \frac{1}{2}(H_z - L_z)\widehat{z} \quad (48)$$

where L and H are the lower and higher bounds on the actual box. During the preprocessing stage, the coordinates of the charges are transformed into the $[-1:1]$ range using the inverse of the above transform, and the matrices of Eqns 44 and 45 are computed. The scaled matrix

$$\mathbf{K}_F = \mathbf{J}^{-1} \cdot \widehat{\mathbf{K}}^F \quad (49)$$

is computed, where $\mathbf{J}$ is the Jacobian of the transformation. Due to transformation properties of the torque, the torque matrix is unchanged by the transformation, as is the energy vector:

$$\mathbf{K}^E = \widehat{\mathbf{K}}^E \quad (50)$$

$$\mathbf{K}^M = \widehat{\mathbf{K}}^M \quad (51)$$

References

1. Alberts, B.; Johnson, A.; Lewis, J.; Raff, M.; Roberts, K.; Walter, P. *Molecular Biology of the Cell*. Fourth Edition. New York: Garland Science; 2002.
2. Pollard, T.; Earnshaw, W. *Cell Biology*. New York: W.B.Saunders; 2007.
3. Elcock, A. *Numerical Computer Methods*, Part D, Vol. 383 of *Methods in Enzymology*. 525 B Street, Suite 1900, San Diego, CA 92101-4495 USA: Academic Press Inc; 2004. Molecular simulations of diffusion and association in multimacromolecular systems; p. 166-198.
4. Ermak D, McCammon J. Brownian Dynamics with Hydrodynamic Interactions. *J. Chem. Phys* 1978;69(4):1352–1360.
5. Jendrejack R, Graham M, de Pablo J. Hydrodynamic interactions in long chain polymers: Application of the Chebyshev polynomial approximation in stochastic simulations. *J. Chem. Phys* 2000;113(7):2894–2900.
6. Shen T, Wong C, McCammon J. Atomistic Brownian dynamics simulation of peptide phosphorylation. *J. Am. Chem. Soc* 2001;123(37):9107–9111. [PubMed: 11552818]
7. Allison S, McCammon J. Dynamics Of Substrate Binding To Copper-Zinc Superoxide-Dismutase. *J. Phys. Chem* 1985;89(7):1072–1074.
8. Madura J, Briggs J, Wade R, Davis M, Luty B, Ilin A, Antosiewicz J, Gilson M, Bagheri B, Scott L, McCammon J. Electrostatics And Diffusion of Molecules in Solution - Simulations With The University-of-Houston Brownian Dynamics Program. *Comput. Phys. Commun* 1995;91(1–3):57–95.

9. Wade R, Luty B, Demchuk E, Madura J, Davis M, Briggs J, McCammon J. Simulation of Enzyme-Substrate encounter with Gated Active Sites. *Nat. Struct. Biol* 1994;1(1):65–69. [PubMed: 7656010]
10. Chang C, Shen T, Trylska J, Tozzini V, McCammon J. Gated binding of ligands to HIV-1 protease: Brownian dynamics simulations in a coarse-grained model. *Biophys. J* 2006;90(11):3880–3885. doi:10.1529/biophysj.105.074575. [PubMed: 16533835]
11. Gabdouliline R, Wade R. Simulation of the diffusional association of Barnase and Barstar. *Biophys. J* 1997;72(5):1917–1929. [PubMed: 9129797]
12. Elcock A, Sept D, McCammon J. Computer simulation of protein-protein interactions. *J. Phys. Chem. B* 2001;105(8):1504–1518. doi:10.1021/jp003602d.
13. Northrup S, Laughner T, Stevenson G. Macrodox macromolecular simulation program. 1997
14. Honig B, Nicholls A. Classical Electrostatics in Biology And Chemistry. *Science* 1995;268(5214):1144–1149. [PubMed: 7761829]
15. Baker N, Sept D, Joseph S, Holst M, McCammon J. Electrostatics of nanosystems: Application to microtubules and the ribosome. *Proc. Natl. Acad. Sci. U. S. A* 2001;98(18):10037–10041. [PubMed: 11517324]
16. Fixman M. Simulation of Polymer Dynamics I. General Theory. *Journal of Chemical Physics* 1978;69(4):1527–1537.
17. Gardiner, CW. *Handbook of Stochastic Methods*. 2nd Edition. Berlin: Springer; 1985.
18. Dickinson E, Allison S, McCammon J. Brownian Dynamics with Rotation Translation Coupling. *Journal of the Chemical Society - Faraday Transactions II* 1985 APR;81:591–601.
19. Allison S. A Brownian Dynamics Algorithm for Arbitrary Rigid Bodies - Application to Polarized Dynamic Light-Scattering. *Macromolecules* 1991;24(2):530–536.
20. Kim, S.; Karila, S. *Microhydrodynamics: Principles and Selected Applications*. Boston: Butterworth-Heinemann; 1991.
21. Fernandes M, Garcia de la Torre J. Brownian dynamics simulation of rigid particles of arbitrary shape in external fields. *Biophysical Journal* 2002;83(6):3039–3048. [PubMed: 12496076]
22. Beard D, Schlick T. Unbiased rotational moves for rigid-body dynamics. *Biophysical Journal* 2003;85(5):2973–2976. [PubMed: 14581199]
23. Levitt M. Simplified Representation of Protein Conformations For Rapid Simulation of Protein Folding. *J. Mol. Biol* 1976;104(1):59–107. [PubMed: 957439]
24. Luty B, McCammon J, Zhou H. Diffusive Reaction-Rates From Brownian Dynamics Simulations - Replacing The Outer Cutoff Surface By An Analytical Treatment. *J. Chem. Phys* 1992;97(8):5682–5686.
25. Huber G, Kim S. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys. J* 1996;70(1):97–110. [PubMed: 8770190]
26. Opendx. 2006. <http://www.opendx.org>
27. Cooper, C. Using Expat. 1999. <http://www.libexpat.org/>
28. Leroy, X.; Doligez, D.; Garrigue, J.; Rémy, D.; Vouillon, J. The Objective Caml System: release 3.11. 2008. <http://caml.inria.fr/pub/docs/misc-ocaml/index.html>
29. I. Free Software Foundation. GNU make manual. 2006. <http://www.gnu.org/software/make/>
30. Barney, B. Posix threads programming. 2009. <https://computing.llnl.gov/tutorials/pthreads/>
31. Humphrey W, Dalke A, Schulten K. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics* 1996;14:33–38. [PubMed: 8744570]
32. Ding H, Karasawa N, Goddard W. Atomic Level Simulations On A Million Particles - The Cell Multipole Method For Coulomb And London Nonbond Interactions. *J. Chem. Phys* 1992;97(6):4309–4315.
33. Gabdouliline R, Wade R. Effective charges for macromolecules in solvent. *J. Phys. Chem* 1996;100(9):3868–3878.
34. Still W, Tempczyk A, Hawley R, Hendrickson T. Semianalytical Treatment of Solvation For Molecular Mechanics And Dynamics. *J. Am. Chem. Soc* 1990;112(16):6127–6129.
35. Visscher PB, Apalkov DM. Simple recursive implementation of fast multipole method. *J. Magn. Magn. Mater* 2010;322(2):275–281. doi:10.1016/j.jmmm.2009.09.033.

36. Rotne J, Prager S. Variational Treatment of Hydrodynamic Interaction in Polymers. *Journal of Chemical Physics* 1969;50(11) 4831–&.
37. Garcia de la Torre J, Bloomfield V. Hydrodynamic Properties of Macromolecular Complexes .1. Translation. *Biopolymers* 1977;16(8):1747–1763.
38. Hansen S. Translational friction coefficients for cylinders of arbitrary axial ratios estimated by Monte Carlo simulation. *J. Chem. Phys* 2004;121(18):9111–9115. doi:10.1063/1.1803533. [PubMed: 15527378]
39. Klosowski J, Held M, Mitchell J, Sowizral H, Zikan K. Efficient collision detection using bounding volume hierarchies of k-DOPS. *IEEE Trans. Vis. Comput. Graph* 1998;4(1):21–36.
40. Welzl E. Smallest Enclosing Disks (Balls And Ellipsoids). *Lect. Notes Comput. Sci* 1991;555:359–370.
41. Goldsmith J, Salmon J. Automatic Creation of Object Hierarchies For Ray Tracing. *IEEE Comput. Graph. Appl* 1987;7(5):14–20.
42. Connolly M. Analytical Molecular-Surface Calculation. *J. Appl. Crystallogr* 1983 Oct;16:548–558.
43. Sanner M, Olson A, Spehner J. Reduced surface: An efficient way to compute molecular surfaces. *Biopolymers* 1996;38(3):305–320. [PubMed: 8906967]
44. Northrup S, Allison S, McCammon J. Brownian Dynamics Simulation of Diffusion-Influenced Bimolecular Reactions. *J. Chem. Phys* 1984;80(4):1517–1526.
45. Politis D, Romano J. The Stationary Bootstrap. *J. Am. Stat. Assoc* 1994;89(428):1303–1313.
46. Lamm G, Schulten K. Extended Brownian Dynamics Approach To Diffusion-Controlled Processes. *J. Chem. Phys* 1981;75(1):365–371.
47. Baerga-Ortiz A, Rezaie A, Komives E. Electrostatic dependence of the thrombin-thrombomodulin interaction. *Journal of Molecular Biology* 2000;296(2):651–658. [PubMed: 10669614]

Figure 1.
Software structure

Figure 2.
Generating trajectories which react or escape